



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

virtual machine and node and monitor

SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used: virtual machine and node and monitor

Found 57,567 of 207,474

Sort results by

relevance

Display results

expanded form

☒ Save results to a Binder☒ Search Tips☐ Open results in a new windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐1 [Devirtualizable virtual machines enabling general, single-node, online maintenance](#)

David E. Lowell, Yasushi Saito, Eileen J. Samberg

 October 2004 **ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 32 , 38 , 39 Issue 5 , 5 , 11

Publisher: ACM Press

Full text available: pdf(174.01 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Maintenance is the dominant source of downtime at high availability sites. Unfortunately, the dominant mechanism for reducing this downtime, cluster rolling upgrade, has two shortcomings that have prevented its broad acceptance. First, cluster-style maintenance over many nodes is typically performed a few nodes at a time, making maintenance slow and often impractical. Second, cluster-style maintenance does not work on single-node systems, despite the fact that their unavailability during maintenance ...

**Keywords:** availability, online maintenance, planned downtime, virtual machines2 [The PDP-11 virtual machine architecture: A case study](#)

Gerald J. Popek, Charles S. Kline

 November 1975 **ACM SIGOPS Operating Systems Review , Proceedings of the fifth ACM symposium on Operating systems principles SOSP '75**, Volume 9 Issue 5

Publisher: ACM Press

Full text available: pdf(905.34 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

At UCLA, a virtual machine system prototype has been constructed for the Digital Equipment Corporation PDP-11/45. In order to successfully implement that system, a number of hardware changes have been necessary. Some overcome basic inadequacies in the original hardware for this purpose, and others enhance the performance of the virtual machine software. Steps in the development of the modified hardware architecture, as well as relevant aspects of the software structure, are discussed. In addition ...

**Keywords:** Computer architecture, Computer security, PDP-11/45, Virtual machine monitor, Virtual machines3 [Formal properties of recursive Virtual Machine architectures.](#)

Gerald Belpaire, Nai-Ting Hsu

 November 1975 **ACM SIGOPS Operating Systems Review , Proceedings of the fifth**

**ACM symposium on Operating systems principles SOSP '75**, Volume 9

Issue 5

**Publisher:** ACM Press

Full text available: pdf(744.44 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A formal model of hardware/software architectures is developed and applied to Virtual Machine Systems. Results are derived on the sufficient conditions that a machine architecture must verify in order to support VM systems. The model deals explicitly with resource mappings (protection) and with I/O devices. Some already published results are retrieved and other ones, more general, are obtained.

**Keywords:** Architecture, Formal requirements, Operating systems, Virtual machine, Virtual machine monitor

**4 Building a World-Wide virtual machine based on web and HPCC technologies**

Kivanc Dincer, Geoffrey C. Fox

November 1996 **Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '96****Publisher:** IEEE Computer Society

Full text available: pdf(229.77 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In today's high performance computing arena, there is a strong trend toward building virtual computers from heterogeneous resources on a network. In this paper we describe our experiences in building a world-wide virtual machine (WWVM) based on emerging Web and existing HPCC technologies. We have constructed a Web-based parallel/distributed programming environment on top of this machine demonstrating MPI and PVM message-passing programs and High Performance Fortran programs. Alternatively, ...

**Keywords:** WWW, HPCC, HPF, MPI, PVM, metacomputing, Perl

**5 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors**

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP '99**, Volume 33 Issue 5**Publisher:** ACM Press

Full text available: pdf(1.93 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a system ...

**6 Testbed directions and experience: PlanetLab: an overlay testbed for broad-coverage services**

Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, Mic Bowman

July 2003 **ACM SIGCOMM Computer Communication Review**, Volume 33 Issue 3**Publisher:** ACM Press

Full text available: pdf(158.92 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

PlanetLab is a global overlay network for developing and accessing broad-coverage network services. Our goal is to grow to 1000 geographically distributed nodes, connected

by a diverse collection of links. PlanetLab allows multiple service to run concurrently and continuously, each in its own slice of PlanetLab. This paper describes our initial implementation of PlanetLab, including the mechanisms used to implement virtualization, and the collection of core services used to manage PlanetLab.

7 Software for high-performance systems: System management software for virtual environments



Geoffroy Vallee, Thomas Naughton, Stephen L. Scott

May 2007 **Proceedings of the 4th international conference on Computing frontiers CF '07**

**Publisher:** ACM Press

Full text available: pdf(293.30 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recently there has been an increased interest in the use of system-level virtualization using mature solutions such as Xen, QEMU, or VMWare. These virtualization platforms are being used in distributed and parallel environments including high performance computing. The use of virtual machines within such environments introduces new challenges to system management. These include tedious tasks such as deploying para-virtualized host operating systems to support virtual machine execution or virtua ...

**Keywords:** OSCAR, configuration management, tools, virtualization

8 Cellular disco: resource management using virtual clusters on shared-memory multiprocessors



Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

**Publisher:** ACM Press

Full text available: pdf(287.05 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a ...

**Keywords:** fault containment, resource management, scalable multiprocessors, virtual machines

9 Work in progress session: A virtual machine monitor for utilizing non-dedicated clusters



Kenji Kaneda, Yoshihiro Oyama, Akinori Yonezawa

October 2005 **Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**

**Publisher:** ACM Press

Full text available: ppt(621.57 KB) Additional Information: [full citation](#)

10 Runtime systems: Proactive fault tolerance for HPC with Xen virtualization



Arun Babu Nagarajan, Frank Mueller, Christian Engelmann, Stephen L. Scott

June 2007 **Proceedings of the 21st annual international conference on Supercomputing ICS '07**

**Publisher:** ACM Press

Full text available: pdf(526.85 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Large-scale parallel computing is relying increasingly on clusters with thousands of

processors. At such large counts of compute nodes, faults are becoming common place. Current techniques to tolerate faults focus on reactive schemes to recover from faults and generally rely on a checkpoint/restart mechanism. Yet, in today's systems, node failures can often be anticipated by detecting a deteriorating health status.

Instead of a reactive scheme for fault tolerance (FT), we are promoting ...

**Keywords:** high-performance computing, proactive fault tolerance, virtualization

11 Novel access control techniques: Chinese-wall process confinement for practical distributed coalitions



Yasuharu Katsuno, Yuji Watanabe, Sanehiro Furuichi, Michiharu Kudo

June 2007 **Proceedings of the 12th ACM symposium on Access control models and technologies SACMAT '07**

**Publisher:** ACM Press

Full text available: pdf(398.48 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A distributed coalition supports distributed mandatory access controls for resources whose security policies differ for each group of components over nodes, and provides secure information operations and exchanges with nodes that handle information over which conflicts of interest may occur. Many projects have proposed distributed coalitions using a virtual machine monitor, but this approach for strong confinement tends to hinder successful deployments in real world scenarios that involve com ...

**Keywords:** distributed coalition, distributed system, mandatory access control, process confinement

12 Distributed VEEs: The entropia virtual machine for desktop grids



Brad Calder, Andrew A. Chien, Ju Wang, Don Yang

June 2005 **Proceedings of the 1st ACM/USENIX international conference on Virtual execution environments VEE '05**

**Publisher:** ACM Press

Full text available: pdf(280.20 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Desktop distributed computing allows companies to exploit the idle cycles on pervasive desktop PC systems to increase the available computing power by orders of magnitude (10x - 1000x). Applications are submitted, distributed, and run on a grid of desktop PCs. Since the applications may be malformed, or malicious, the key challenges for a desktop grid are how to 1) prevent the distributed computing application from unwarranted access or modification of data and files on the desktop PC, 2) contro ...

**Keywords:** desktop grids, grid computing, virtual machine

13 Virtual machine-based simulation of distributed computing and network computing



Richard T. Wang, J. C. Browne

September 1981 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1981 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '81, Volume 10 Issue 3**

**Publisher:** ACM Press

Full text available: pdf(255.73 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper proposes the use of virtual machine architectures as a means of modeling and analyzing networks and distributed computing systems. The requirements for such modeling and analysis are explored and defined along with an illustrative study of an X.25 link-level protocol performance under normal execution conditions. The virtualizable architecture used in this work is the Data General Nova 3/D.

14 Work in progress session: A virtual machine monitor for utilizing non-dedicated clusters



Kenji Kaneda, Yoshihiro Oyama, Akinori Yonezawa

October 2005 **Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**

**Publisher:** ACM Press

Full text available: pdf(383.72 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We have designed and implemented a virtual machine monitor (VMM) for utilizing non-dedicated clusters. The VMM virtualizes a shared-memory multi-processor machine on a commodity cluster. In addition, it hides dynamic changes of physical hardware configurations. The experimental result demonstrates the feasibility of our approach.

**Keywords:** distributed systems, single system image, virtual machine monitors

15 Cluster resource management: An integrated experimental environment for distributed systems and networks



Brian White, Jay Lepreau, Leigh Stoller, Robert Ricci, Shashi Guruprasad, Mac Newbold, Mike Hibler, Chad Barb, Abhijeet Joglekar

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue S1

**Publisher:** ACM Press

Full text available: pdf(2.10 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Three experimental environments traditionally support network and distributed systems research: network emulators, network simulators, and live networks. The continued use of multiple approaches highlights both the value and inadequacy of each. Netbed, a descendant of Emulab, provides an experimentation facility that integrates these approaches, allowing researchers to configure and access networks composed of emulated, simulated, and wide-area nodes and links. Netbed's primary goals are ease ...

16 SHARP: an architecture for secure resource peering



Yun Fu, Jeffrey Chase, Brent Chun, Stephen Schwab, Amin Vahdat

October 2003 **ACM SIGOPS Operating Systems Review , Proceedings of the nineteenth ACM symposium on Operating systems principles SOSP '03**, Volume 37 Issue 5

**Publisher:** ACM Press

Full text available: pdf(339.51 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents Sharp, a framework for secure distributed resource management in an Internet-scale computing infrastructure. The cornerstone of Sharp is a construct to represent cryptographically protected resource <it>claims</it>---promises or rights to control resources for designated time intervals---together with secure mechanisms to subdivide and delegate claims across a network of resource managers. These mechanisms enable flexible <it>resource peering ...

**Keywords:** peer-to-peer, resource allocation, resource peering

17 Disco: running commodity operating systems on scalable multiprocessors



Edouard Bugnion, Scott Devine, Mendel Rosenblum

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

**Publisher:** ACM Press

Full text available: pdf(2.30 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

18 Superconcurrent simulation of polymer chains on heterogeneous networks

H. Nakanishi, V. Rego, V. Sunderam

December 1992 **Proceedings of the 1992 ACM/IEEE conference on Supercomputing Supercomputing '92**

**Publisher:** IEEE Computer Society Press


Full text available:  pdf(792.54 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

19 Performance modeling for the panda array I/O library

Ying Chen, Marianne Winslett, Szu-wen Kuo, Yong Cho, Mahesh Subramaniam, Kent Seamons

November 1996 **Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '96**

**Publisher:** IEEE Computer Society

Full text available:  pdf(244.54 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present an analytical performance model for Panda, a library for synchronized i/o of large multidimensional arrays on parallel and sequential platforms, and show how the Panda developers use this model to evaluate Panda's parallel i/o performance and guide future Panda development. The model validation shows that system developers can simplify performance analysis, identify potential performance bottlenecks, and study the design trade-offs for Panda on massively parallel platforms more e ...


20 Disco: running commodity operating systems on scalable multiprocessors



Edouard Bugnion, Scott Devine, Kinshuk Govil, Mendel Rosenblum

November 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 4

**Publisher:** ACM Press

Full text available:  pdf(400.76 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In this article we examine the problem of extending modern operating systems to run efficiently on large-scale shared-memory multiprocessors without a large implementation effort. Our approach brings back an idea popular in the 1970s: virtual machine monitors. We use virtual machines to run multiple commodity operating systems on a scalable multiprocessor. This solution addresses many of the challenges facing the system software for these machines. We demonstrate our approach with a prototy ...

**Keywords:** scalable multiprocessors, virtual machines

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)